arXiv:1610.09887v1 [cs.LG] 31 Oct 2016

Depth Separation in ReLU Networks for Approximating Smooth Non-Linear Functions

Itay Safran
Weizmann Institute of Science
itay.safran@weizmann.ac.il

Ohad Shamir
Weizmann Institute of Science
ohad.shamir@weizmann.ac.il

Abstract

We provide a depth-based separation result for feed-forward ReLU neural networks, showing that a wide family of non-linear, twice-differentiable functions on $[0,1]^d$, which can be approximated to accuracy ϵ by ReLU networks of depth and width $\mathcal{O}(\text{poly}\,(\log(1/\epsilon)))$, cannot be approximated to similar accuracy by constant-depth ReLU networks, unless their width is at least $\Omega(1/\epsilon)$.

1 Introduction

Deep learning, in the form of artificial neural networks have seen a dramatic resurgence in the past recent years, achieving great improvement over various fields of artificial intelligence such as computer vision and speech recognition. While empirically successful, our theoretical understanding of them is still limited at best.

In this work, we focus on the aspect of the *expressive power* of neural networks. Specifically, We investigate how the accuracy to which they are able to approximate various target functions increases with their size, and in particular, what is the trade-off between increasing their width vs. increasing their depth. It is well-known that networks of depth 2 can already approximate any continuous target function on the boolean hypercube $[0,1]^d$ to arbitrary accuracy, albeit at the cost of width exponential in the dimension d [Cybenko, 1989]. In contrast, it has long been evident that deeper networks tend to perform better than shallow ones, a phenomenon supported by the intuition that depth, providing compositional expressibility, is necessary for efficiently representing some functions. Indeed, recent empirical evidence suggests that even at large depths, deeper networks can offer benefits over shallower networks [He et al., 2015].

In this paper, we consider feed-forward neural networks with the commonly-used rectified linear unit (ReLU, $[z]_+ = \max\{0,z\}$) as the activation function. We focus mainly on the following question: *Given some function $f : [0,1]^d \to \mathbb{R}$ and accuracy parameter ϵ, what lower and upper bounds can be derived on the required width and depth of a ReLU neural network, in order to approximate f up to an accuracy of ϵ?* Here, we define accuracy in terms of expected squared loss, or sometimes maximal loss, over some domain (see Sec. 2 for more details).

In the first part of our work (Sec. 3), we discuss lower bounds on the size of a network required for achieving approximation error of ϵ for functions which are C^2 (that is, twice-continuously differentiable) on the bounded domain $[0,1]^d$. The approximation quality is measured in terms of the expected squared loss w.r.t. the uniform distribution on the domain. Essentially, we show that *any* $f : [0,1]^d \to \mathbb{R}$ in C^2, which is "non-linear" in a portion of its domain (in terms of having non-zero curvature, at least in some direction),

cannot be approximated to accuracy ϵ (for arbitrary $\epsilon > 0$) using ReLU networks of constant depth, unless the width is at least of order $1/\epsilon$.

In the second part of our work (Sec. 4), we show that a wide class of functions on the domain $[0,1]^d$ (including some of those for which our lower bounds apply) can be approximated to accuracy ϵ using a network of width and depth only $\mathcal{O}\left(\text{poly}\left(\log\left(\frac{1}{\epsilon}\right)\right)\right)$. Specifically, this is shown for functions which (i) Can be approximated well using a series of at most $\mathcal{O}\left(\text{poly}\left(\log\left(\frac{1}{\epsilon}\right)\right)\right)$ operations involving additions and multiplications; and (ii) Each intermediate stage in the series of computations has its value bounded by a quantity not larger than $\mathcal{O}\left(\text{poly}\left(\frac{1}{\epsilon}\right)\right)$.

In Sec. 5, we combine the upper and lower bounds, to derive an exponential separation result (in the accuracy parameter ϵ) for the wide class of functions lying in the intersection of those described earlier: That is, for any non-linear C^2 function, which can be efficiently approximated using additions and multiplications.

Related Work

The question of studying the effect of depth in neural networks has received considerable attention recently. Here, we briefly review some of the most relevant work, providing formal depth separation results.

In Eldan and Shamir [2016], the authors show that depth in neural networks can be exponentially valuable even if increased only by 1, by proving the existence of a family of radial functions which are expressible using a depth 3 feed-forward neural network of size polynomial in the dimension d, whereas a network of depth 2 cannot approximate these functions to more than constant accuracy (even if the activation function is arbitrary), unless its width is exponential in d. However, it is not clear whether their proof techniques extend to deeper networks. In contrast, our results apply to potentially deeper networks, and the separation is in terms of the required accuracy ϵ, rather than some fixed depths.

The work perhaps most similar to ours is Telgarsky [2016]. In that work, the author proves a separation result between networks of depth k and depth $o\left(\frac{k}{\log(k)}\right)$ (for arbitrary k), by showing that in one dimension, a depth k network can realize a saw-tooth function with $\exp(\mathcal{O}(k))$ oscillations, whereas any network of depth $o\left(\frac{k}{\log(k)}\right)$ would require a width super-polynomial in k to approximate it by more than a constant. Like our work, it potentially applies to networks of arbitrary depth, and the trade-off between depth and width is roughly the same (doubling the depth is equivalent to squaring the width). However, the papers differ in two main aspects: (i) The notion of separation is in a parameter k rather than the accuracy to which the target function is approximated to; (ii) Unlike the tailored construction of a highly oscillatory function with Lipschitz constant exponential in k, our result applies to a wide variety of more naturally encountered functions, including functions with a bounded Lipschitz constant.

Finally, in Cohen et al. [2016], the authors establish a depth separation result for networks with a certain tensor structure. The authors achieve a very strong separation; not only does it show the existence of target functions expressible by deep networks which are hard to approximate using shallow networks, it in fact shows that the class of functions expressible by shallower networks has measure zero among the class of functions expressible using deep networks. Albeit very powerful, this separation result does not apply for the kind of ReLU neural networks discussed in this paper.

2 Preliminaries

To formalize our setup, we recall that a fully-connected feed-forward artificial neural network computes a function $\mathbb{R}^d \to \mathbb{R}$, and is composed of neurons connected according to a directed acyclic graph. Specifically,

the neurons can be decomposed into layers, where the output of each neuron is connected to all neurons in the succeeding layer and them alone. We focus on ReLU networks, where each neuron computes a function of the form $\mathbf{x} \mapsto \left[\mathbf{w}^\top \mathbf{x} + b\right]_+$ where $\mathbf{w}$ is a weight vector, b is a bias term specific to that neuron and $[z]_+ = \max\{0, z\}$ is the ReLU activation function.

For a vector $\mathbf{b} = (b_1, \ldots, b_n)$ and a matrix $W = (\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_n)^\top$, we let $[W\mathbf{x} + \mathbf{b}]_+$ be a shorthand for $\left(\left[\mathbf{w}_1^\top \mathbf{x} + b_1\right]_+, \ldots, \left[\mathbf{w}_n^\top \mathbf{x} + b_n\right]_+\right)$. We define a layer of n neurons as

$$\mathbf{x} \mapsto [W\mathbf{x} + \mathbf{b}]_+ .$$

Finally, by denoting the output of the i^{th} layer as O_i, we can define a network of arbitrary depth recursively by

$$O_{i+1} = [W_{i+1} O_i + \mathbf{b}_{i+1}]_+ ,$$

where $W_i, \mathbf{b}_i$ represent the matrix of weights and bias of the i^{th} layer, respectively. Following a standard convention for multi-layer networks, the final layer h is a purely linear function with no bias, i.e.

$$O_h = W_h \cdot O_{h-1}.$$

Define the *depth* of the network as the number of layers l, and denote the number of neurons n_i in the i^{th} layer as the *size* of the layer. We define the *width* of a network as $\max_{i \in \{1,\ldots,l\}} n_i$.

Given some function f on $[0,1]^d$ and approximation $\tilde{f}$, we generally define the quality of approximation as

$$\int_{[0,1]^d} (f - \tilde{f})^2 d\mu_d,$$

where μ_d is the uniform measure on $[0,1]^d$. This corresponds to the expected squared loss w.r.t. the uniform distribution. We refer to this as approximation in the L_2-norm sense. In some of the results, we also consider approximation with respect to the infinity norm, defined as

$$\sup_{\mathbf{x} \in [0,1]^d} |f(\mathbf{x}) - \tilde{f}(\mathbf{x})|.$$

Clearly, this upper-bounds the (square root of the) L_2 approximation error defined above.

3 Lower Bounds on the Accuracy of Approximating C^2 Functions

In this section, we establish a lower bound for approximating C^2 functions using ReLU neural networks. We first note that such a bound can be trivial if we consider cases where the function being approximated is piece-wise linear and continuous, as in this case, the function might be possible to approximate to accuracy 0 using a large enough network. To ensure the lower bound we derive considers these trivial cases, we need to introduce a measure of how 'non-linear' our target function is. This non-linearity of the target function is captured by the following quantity:

Definition 1. *Let μ_d denote the uniform measure on $[0,1]^d$. For a function $f : [0,1]^d \to \mathbb{R}$ and some $\lambda > 0$, denote*

$$\sigma_\lambda(f) = \sup_{U \in \mathcal{U} \text{ s.t. } \|H(f)(\mathbf{x})\|_{sp} \geq \lambda \ \forall \mathbf{x} \in U} \mu_d(U),$$

where $\mathcal{U}$ is the set of all connected and measurable subsets of $[0,1]^d$, and $\|H(f)(\mathbf{x})\|_{sp}$ is the spectral norm of the Hessian of f at $\mathbf{x}$.

In words, $\sigma_\lambda(f)$ is the measure (w.r.t. the uniform distribution on $[0,1]^d$) of the largest connected set in the domain of f, where at any point, f has curvature at least λ along some direction. Clearly, if f is piece-wise linear then $\lambda = 0$ for all $\lambda > 0$. We note that whenever clear from context, we omit f from the notation and use σ_λ instead.

We can now state the main result of this section:

Theorem 1. *Let $\mathcal{G}_n$ be the family of piece-wise linear functions on the domain $[0,1]$ comprised of at most n linear segments. Let $\mathcal{G}_n^d$ be the family of piece-wise linear functions on the domain $[0,1]^d$, with the property that for any $g \in \mathcal{G}_n^d$ and any affine transformation $h : \mathbb{R} \to \mathbb{R}^d$, $g \circ h \in \mathcal{G}_n$. Suppose $f : [0,1]^d \to \mathbb{R}$ is C^2. Then for all $\lambda > 0$*

$$\inf_{g \in \mathcal{G}_n^d} \int_{[0,1]^d} (f-g)^2 d\mu_d \geq \frac{c \cdot \lambda^2 \cdot \sigma_\lambda(f)^5}{n^4},$$

where $c > 0$ is a universal constant.

Thm. 1 establishes that the error of a piece-wise linear approximation of a C^2 function cannot decay faster than quartically in the number of linear segments of any *one-dimensional* projection of the approximating function. Note that this result is stronger than a bound in terms of the total number of linear regions in $\mathbb{R}^d$, since that number can be exponentially higher (in the dimension) than n as defined in the theorem.

To translate this result to the context of ReLU neural networks, we use the result in Telgarsky [2016, Lemma 3.2], of which the following is an immediate corollary.

Corollary 1. *Let $\mathcal{N}_{m,l}^d$ denote the family of ReLU neural networks receiving input of dimension d and having depth l and maximal width m. Then*

$$\mathcal{N}_{m,l}^d \subseteq \mathcal{G}_{(2m)^l}^d.$$

The following corollary bears great resemblance to the lower bound provided in Telgarsky [2016]. Albeit lower bounding the accuracy to which a ReLU network of a given size can approximate a certain function, rather than indicating what is a lower bound on the minimal size required for achieving non-constant approximation error. By combining Thm. 1 and Corollary 1, we have that

Corollary 2. *Suppose $f : [0,1]^d \to \mathbb{R}$ is C^2. Then for all $\lambda > 0$,*

$$\inf_{g \in \mathcal{N}_{m,l}^d} \int_{[0,1]^d} (f-g)^2 d\mu_d \geq \frac{c \cdot \lambda^2 \cdot \sigma_\lambda^5}{2^l m^l}.$$

Corollary 2 conveys the key tradeoff between depth and width when approximating C^2 functions using ReLU networks: The error cannot decay faster than polynomially in the width, yet the bound deteriorates exponentially in the depth. As we show in the following section, this deterioration does not stem from the bound in Corollary 2 being loose: For well-behaved f, it is indeed possible to construct ReLU networks, where the approximation error decays exponentially with depth.

In the rest of the section, we prove Thm. 1 by a series of intermediate results, some of which may be of independent interest. In a nutshell, we show that strictly curved functions cannot be well-approximated by piecewise linear functions, unless the number of regions is large. To that end, we first establish some necessary tools based on Legendre polynomials. We then prove a result specific to the one-dimensional case, including an explicit lower bound if the target function is quadratic (Thm. 3) or strongly convex or concave (Thm. 4). Finally, we expand the construction to get a result in general dimension d.

3.1 Some Technical Tools

Definition 2. *Let P_i denote the i^{th} Legendre Polynomial given by Rodrigues' formula:*

$$P_i(x) = \frac{1}{2^i i!}\frac{d^i}{dx^i}\left[\left(x^2-1\right)^i\right].$$

These polynomials are useful for the following analysis since they obey the orthogonality relationship

$$\int_{-1}^{1} P_i(x)\,P_j(x)\,dx = \frac{2}{2i+1}\delta_{ij}.$$

Since we are interested in approximations on small intervals where the approximating function is linear, we use the change of variables $x = \frac{2}{\ell}t - \frac{2}{\ell}a - 1$ to obtain an orthogonal family $\left\{\tilde{P}_i\right\}_{i=1}^{\infty}$ of shifted Legendre polynomials on the interval $[a, a+\ell]$ with respect to the L_2 norm. The first few polynomials of this family are given by

$$\begin{aligned}
\tilde{P}_0(x) &= 1\\
\tilde{P}_1(x) &= \frac{2}{\ell}x - \left(\frac{2}{\ell}a + 1\right)\\
\tilde{P}_2(x) &= \frac{6}{\ell^2}x^2 - \left(\frac{12a}{\ell^2} + \frac{6}{\ell}\right)x + \left(\frac{6a^2}{\ell^2} + \frac{6a}{\ell} + 1\right).
\end{aligned} \tag{1}$$

The shifted Legendre polynomial obey the orthogonality relationship

$$\int_{a}^{a+\ell} \tilde{P}_i(x)\,\tilde{P}_j(x)\,dx = \frac{\ell}{2i+1}\delta_{ij}. \tag{2}$$

Definition 3. *We define the Fourier-Legendre series of a function $f : [a, a+\ell] \to \mathbb{R}$ to be*

$$f(x) = \sum_{i=0}^{\infty} \tilde{a}_i \tilde{P}_i(x),$$

where the Fourier-Legendre Coefficients $\tilde{a}_i$ are given by

$$\tilde{a}_i = \frac{2i+1}{\ell}\int_{a}^{a+\ell} \tilde{P}_i(x)\,f(x)\,dx.$$

Theorem 2. *A generalization of Parseval's identity yields*

$$\|f\|_{L_2}^2 = \ell\sum_{i=0}^{\infty}\frac{\tilde{a}_i^2}{2i+1}. \tag{3}$$

Definition 4. *A function f is* λ-strongly convex *if for all $\mathbf{w}, \mathbf{u}$ and $\alpha \in (0,1)$,*

$$f(\alpha\mathbf{w} + (1-\alpha)\mathbf{u}) \le \alpha f(\mathbf{w}) + (1-\alpha)f(\mathbf{u}) - \frac{\lambda}{2}\alpha(1-\alpha)\|\mathbf{w}-\mathbf{u}\|_2^2.$$

A function is λ-strongly concave*, if $-f$ is λ-strongly convex.*

3.2 One-dimensional Lower Bounds

We begin by proving two useful lemmas; the first will allow us to compute the error of a linear approximation of one-dimensional functions on arbitrary intervals, and the second will allow us to infer bounds on the entire domain of approximation, from the lower bounds we have on small intervals where the approximating function is linear.

Lemma 1. *Let $f \in C^2$. Then the error of the optimal linear approximation of f denoted Pf on the interval $[a, a+\ell]$ satisfies*

$$\|f - Pf\|_{L_2}^2 = \ell \sum_{i=2}^{\infty} \frac{\tilde{a}_i^2}{2i+1}. \tag{4}$$

Proof. A standard result on Legendre polynomials is that given any function f on the interval $[a, a+\ell]$, the best linear approximation (w.r.t. the L_2 norm) is given by

$$Pf = \tilde{a}_0 \tilde{P}_0(x) + \tilde{a}_1 \tilde{P}_1(x),$$

where $\tilde{P}_0, \tilde{P}_1$ are the shifted Legendre polynomials of degree 0 and 1 respectively, and $\tilde{a}_0, \tilde{a}_1$ are the first two Fourier-Legendre coefficients of f as defined in Eq. (3). The square of the error obtained by this approximation is therefore

$$\begin{aligned}
\|f - Pf\|^2 &= \|f\|^2 - 2\langle f, Pf\rangle + \|Pf\|^2 \\
&= \ell \left(\sum_{i=0}^{\infty} \frac{\tilde{a}_i^2}{2i+1} - 2\left(\tilde{a}_0^2 + \frac{\tilde{a}_1^2}{3}\right) + \tilde{a}_0^2 + \frac{\tilde{a}_1^2}{3} \right) \\
&= \ell \sum_{i=2}^{\infty} \frac{\tilde{a}_i^2}{2i+1}.
\end{aligned}$$

Where in the second equality we used the orthogonality relationship from Eq. (2), and the generalized Parseval's identity from Eq. (3). □

Lemma 2. *Suppose $f : [0,1] \to \mathbb{R}$ satisfies $\|f - Pf\|_{L_2}^2 \geq c\ell^5$ for some constant $c > 0$, and on any interval $[a, a+\ell] \subseteq [0,1]$. Then*

$$\inf_{g \in \mathcal{G}_n} \int_0^1 (f-g)^2 d\mu \geq \frac{c}{n^4}.$$

Proof. Let $g \in \mathcal{G}_n$ be some function, let $a_0 = 0, a_1, \ldots, a_{n-1}, a_n = 1$ denote its partition into segments of length $\ell_j = a_j - a_{j-1}$, where g is linear when restricted to any interval $[a_{j-1}, a_j]$, and let $g_j,\ j = 1, \ldots, n$ denote the linear restriction of g to the interval $[a_{j-1}, a_j]$. Then

$$\begin{aligned}
\int_0^1 (p-g)^2 \, d\mu &= \sum_{j=1}^{n} \int_{a_{j-1}}^{a_j} (p - g_j)^2 \, d\mu \\
&\geq \sum_{j=1}^{n} c\ell_j^5 \\
&= c \sum_{j=1}^{n} \ell_j^5.
\end{aligned} \tag{5}$$

Now, recall Hölder's sum inequality which states that for any p, q satisfying $\frac{1}{p} + \frac{1}{q} = 1$ we have

$$\sum_{j=1}^{n} |x_j y_j| \leq \left(\sum_{j=1}^{n} |x_j|^p \right)^{1/p} \left(\sum_{j=1}^{n} |y_j|^q \right)^{1/q}.$$

Plugging in $x_j = \ell_j, y_j = 1 \ \ \forall j \in \{1, \ldots, n\}$ we have

$$\sum_{j=1}^{n} |\ell_j| \leq \left(\sum_{j=1}^{n} |\ell_j|^p \right)^{1/p} n^{1/q},$$

and using the equalities $\sum_{j=1}^{n} |\ell_j| = 1$ and $\frac{p}{q} = p - 1$ we get that

$$\frac{1}{n^{p-1}} \leq \sum_{j=1}^{n} |\ell_j|^p . \tag{6}$$

Plugging the inequality from Eq. (6) with $p = 5$ in Eq. (5) yields

$$\int_0^1 (p - g)^2 \, d\mu \geq \frac{c}{n^4},$$

concluding the proof of the lemma. □

Our first lower bound for approximation using piece-wise linear functions is for non-linear target functions of the simplest kind. Namely, we obtain lower bounds on quadratic functions.

Theorem 3. *If $\mathcal{G}_n$ is the family of piece-wise linear functions with at most n linear segments in the interval $[0, 1]$, then for any quadratic function $p(x) = p_2 x^2 + p_1 x + p_0$, we have*

$$\inf_{g \in \mathcal{G}_n} \int_0^1 (p - g)^2 d\mu \geq \frac{p_2^2}{180 n^4}. \tag{7}$$

Proof. Observe that since p is a degree 2 polynomial, we have that its coefficients satisfy $\tilde{a}_i = 0 \ \forall i \geq 3$, so from Lemma 1 its optimal approximation error equals $\frac{\tilde{a}_2^2 \ell}{5}$. Computing $\tilde{a}_2$ can be done directly from the equation

$$p(x) = \sum_{i=0}^{2} \tilde{a}_i \tilde{P}_i(x),$$

Which gives $\tilde{a}_2 = \frac{p_2 \ell^2}{6}$ due to Eq. (1). This implies that

$$\|p - Pp\|^2 = \frac{p_2^2 \ell^5}{180}.$$

Note that for quadratic functions, the optimal error is dependent solely on the length of the interval. Using Lemma 2 with $c = \frac{p_2^2}{180}$ we get

$$\int_0^1 (p - g)^2 \, d\mu \geq \frac{p_2^2}{180 n^4},$$

concluding the proof of the theorem. □

Computing a lower bound for quadratic functions is made easy since the bound on any interval $[a, a+\ell]$ depends on ℓ but not on a. This is not the case in general, as can be seen by observing monomials of high degree k. As k grows, x^k on the interval $[0, 0.5]$ converges rapidly to 0, whereas on $\left[1-\frac{1}{k}, 1\right]$ its second derivative is lower bounded by $\frac{k(k-1)}{4}$, which indicates that indeed a lower bound for x^k will depend on a.

For non-quadratic functions, however, we now show that a lower bound can be derived under the assumption of strong convexity (or strong concavity) in $[0, 1]$.

Theorem 4. *Suppose $f : [0, 1] \to \mathbb{R}$ is C^2 and either λ-strongly convex or λ-strongly concave. Then*

$$\inf_{g \in \mathcal{G}_n} \int_0^1 (f-g)^2 d\mu \geq c\lambda^2 n^{-4}, \tag{8}$$

where $c > 0$ is a universal constant.

Proof. We first stress that an analogous assumption to λ-strong convexity would be that f is λ-strongly concave, since the same bound can be derived under concavity by simply applying the theorem to the additive inverse of f, and observing that the additive inverse of any piece-wise linear approximation of f is in itself, of course, a piece-wise linear function. For this reason from now on we shall use the convexity assumption, but will also refer without loss of generality to concave functions.

As in the previous proof, we first prove a bound on intervals of length ℓ and then generalize for the unit interval. From Lemma 1, it suffices that we lower bound $\tilde{a}_2$ (although this might not give the tightest lower bound in terms of constants, it is possible to show that it does give a tight bound over all C^2 functions). We compute

$$\begin{aligned}\tilde{a}_2 &= \frac{5}{\ell}\int_a^{a+\ell} \tilde{P}_2(x) f(x)\, dx \\ &= \frac{5}{\ell}\int_a^{a+\ell} P_2\left(\frac{2}{\ell}x - \frac{2}{\ell}a - 1\right) f(x)\, dx,\end{aligned}$$

using the change of variables $t = \frac{2}{\ell}x - \frac{2}{\ell}a - 1$, $dt = \frac{2}{\ell}dx$, we get the above equals

$$\begin{aligned}&\frac{5}{2}\int_{-1}^{1} P_2(t) f\left(\frac{\ell}{2}t + \frac{\ell}{2} + a\right) dt \\ =&\frac{5}{4}\int_{-1}^{1} \left(3t^2-1\right) f\left(\frac{\ell}{2}t + \frac{\ell}{2} + a\right) dt.\end{aligned}$$

We now integrate by parts twice, taking the anti-derivative of the polynomial to obtain

$$\begin{aligned}
&\frac{5}{4}\int_{-1}^{1}\left(3t^2-1\right)f\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\\
=&\frac{5}{4}\left[\left(t^3-t\right)f\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)\right]_{-1}^{1}-\frac{5\ell}{8}\int_{-1}^{1}\left(t^3-t\right)f'\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\\
=&\frac{5\ell}{8}\int_{-1}^{1}\left(t-t^3\right)f'\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\\
=&\frac{5\ell}{8}\left[\left(\frac{t^2}{2}-\frac{t^4}{4}\right)f'\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)\right]_{-1}^{1}\\
&-\frac{5\ell^2}{16}\int_{-1}^{1}\left(\frac{t^2}{2}-\frac{t^4}{4}\right)f''\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\\
=&\frac{5\ell}{32}\left(f'\left(a+\ell\right)-f'\left(a\right)\right)-\frac{5\ell^2}{16}\int_{-1}^{1}\left(\frac{t^2}{2}-\frac{t^4}{4}\right)f''\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt. \qquad (9)
\end{aligned}$$

But since $\frac{t^2}{2}-\frac{t^4}{4}\in\left[0,\frac{1}{4}\right]\ \forall t\in[-1,1]$ and since $f''>0$ due to strong convexity, we have that

$$\int_{-1}^{1}\left(\frac{t^2}{2}-\frac{t^4}{4}\right)f''\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\le\frac{1}{4}\int_{-1}^{1}f''\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt.$$

Plugging this inequality in Eq. (9) yields

$$\begin{aligned}
\tilde{a}_2&\ge\frac{5\ell}{32}\left(f'\left(a+\ell\right)-f'\left(a\right)\right)-\frac{5\ell^2}{64}\int_{-1}^{1}f''\left(\frac{\ell}{2}t+\frac{\ell}{2}+a\right)dt\\
&=\frac{5\ell}{32}\left(f'\left(a+\ell\right)-f'\left(a\right)\right)-\frac{5\ell^2}{64}\left(f'\left(a+\ell\right)-f'\left(a\right)\right)\\
&=\left(1-\frac{\ell}{2}\right)\frac{5\ell}{32}\left(f'\left(a+\ell\right)-f'\left(a\right)\right),
\end{aligned}$$

but $\ell\le 1$, so the above is at least

$$\frac{5\ell}{64}\left(f'\left(a+\ell\right)-f'\left(a\right)\right). \qquad (10)$$

By Lagrange's intermediate value theorem, there exists some $\xi\in[a,a+\ell]$ such that $f'\left(a+\ell\right)-f'\left(a\right)=\ell f''\left(\xi\right)$, so Eq. (10) is at least

$$\frac{5\ell^2}{64}f''\left(\xi\right),$$

and by using the strong convexity of f again, we get that

$$\tilde{a}_2\ge\frac{5\lambda\ell^2}{64}.$$

Lemma 1 now gives

$$\begin{aligned}
\left\|f-Pf\right\|^2&=\ell\sum_{i=2}^{\infty}\frac{\tilde{a}_i^2}{2i+1}\\
&\ge\ell\frac{\tilde{a}_2^2}{5}\\
&\ge\frac{5\lambda^2\ell^5}{4096}.
\end{aligned}$$

Finally, by using Lemma 2 we conclude

$$\inf_{g\in\mathcal{G}_n}\int_0^1 (f-g)^2 d\mu \geq \frac{5\lambda^2}{4096n^4}.$$

□

We now derive a general lower bound for functions $f:[0,1]\to\mathbb{R}$.

Theorem 5. *Suppose $f:[0,1]\to\mathbb{R}$ is piece-wise C^2. Then for any $\lambda>0$*

$$\inf_{g\in\mathcal{G}_n}\int_0^1 (f-g)^2 d\mu \geq \frac{c\cdot\lambda^2\cdot\sigma_\lambda(f)^5}{n^4}.$$

Proof. First, observe that if f is λ-strongly convex on $[a,b]$, then $f\left((b-a)x+a\right)$ is $\lambda(b-a)^2$-strongly convex on $[0,1]$ since $\forall x\in[0,1]$,

$$\frac{\partial}{\partial x^2} f\left((b-a)x+a\right) = (b-a)^2 f''\left((b-a)x+a\right) \geq \lambda(b-a)^2.$$

Now, we use the change of variables $x=(b-a)t+a$, $dx=(b-a)\,dt$

$$\begin{aligned}
&\inf_{g\in\mathcal{G}_n}\int_a^b (f(x)-g(x))^2 dx \\
&= \inf_{g\in\mathcal{G}_n}(b-a)\int_0^1 \left(f\left((b-a)t+a\right)-g\left((b-a)t+a\right)\right)^2 dt \\
&= \inf_{g\in\mathcal{G}_n}(b-a)\int_0^1 \left(f\left((b-a)t+a\right)-g\left(t\right)\right)^2 dt \\
&\geq \frac{c\cdot\lambda^2\cdot(b-a)^5}{n^4}, \qquad (11)
\end{aligned}$$

where the inequality follows from an application of Thm. 4. Back to the theorem statement, if $\sigma_\lambda=0$ then the bound trivially holds, therefore assume $\lambda>0$ such that $\sigma_\lambda>0$. Since f is strongly convex on a set of measure $\sigma_\lambda>0$, the theorem follows by applying the inequality from Eq. (11). □

3.3 Multi-dimensional Lower Bounds

We now move to generalize the bounds in the previous subsection to general dimension d. Namely, we can now turn to proving Thm. 1.

Proof of Thm. 1. Analogously to the proof of Thm. 5, we identify a neighborhood of f in which the restriction of f to a line in a certain direction is non-linear. We then integrate along all lines in that direction and use the result of Thm. 5 to establish the lower bound.

Before we can prove the theorem, we need to assert that indeed there exists a set having a strictly positive measure where f has strong curvature along a certain direction. Assuming f is not piece-wise linear; namely, we have some $\mathbf{x}_0\in[0,1]^d$ such that $H(f)(\mathbf{x}_0)\neq 0$. Since $H(f)$ is continuous, we have that the function $h_{\mathbf{v}}(\mathbf{x})=\mathbf{v}^\top H(f)(\mathbf{x})\,\mathbf{v}$ is continuous and there exists a direction $\mathbf{v}\in\mathbb{R}^d$ where without loss of generality $h_{\mathbf{v}}(\mathbf{x}_0)>0$. Thus, we have an open neighborhood containing $\mathbf{x}_0$ where restricting f to the direction $\mathbf{v}$ forms a strongly convex function, which implies that indeed $\sigma_\lambda>0$ for small enough $\lambda>0$.

We now integrate the approximation error on f in the neighborhood U along the direction $\mathbf{v}$, $\|\mathbf{v}\|_2 = 1$, where $\mathbf{v}$ is the eigenvector of λ, to derive a lower bound. We compute

$$\begin{aligned}
&\inf_{g\in\mathcal{G}_n^d}\int_{[0,1]^d}(f-g)^2 d\mu_d \\
=&\inf_{g\in\mathcal{G}_n^d}\int_{\mathbf{u}:\langle\mathbf{u},\mathbf{v}\rangle=0}\int_{\beta:(\mathbf{u}+\beta\mathbf{v})\in[0,1]^d}(f-g)^2\, d\mu_1 d\mu_{d-1} \\
\geq&\inf_{g\in\mathcal{G}_n^d}\int_{\mathbf{u}:\langle\mathbf{u},\mathbf{v}\rangle=0}\int_{\beta:(\mathbf{u}+\beta\mathbf{v})\in U}(f-g)^2\, d\mu_1 d\mu_{d-1} \\
\geq&\int_{\mathbf{u}:\langle\mathbf{u},\mathbf{v}\rangle=0}\left(\mu_1\left(\{\beta:(\mathbf{u}+\beta\mathbf{v})\in U\}\right)\right)^5\frac{5\lambda^2}{4096n^4}d\mu_{d-1} \\
=&\frac{5\lambda^2}{4096n^4}\int_{\mathbf{u}:\langle\mathbf{u},\mathbf{v}\rangle=0}\left|\mu_1\left(\{\beta:(\mathbf{u}+\beta\mathbf{v})\in U\}\right)\right|^5 d\mu_{d-1} \\
\geq&\frac{5\lambda^2}{4096n^4}\left(\int_{\mathbf{u}:\langle\mathbf{u},\mathbf{v}\rangle=0}\mu_1\left(\{\beta:(\mathbf{u}+\beta\mathbf{v})\in U\}\right)d\mu_{d-1}\right)^5 \\
=&\frac{5\lambda^2\sigma_\lambda^5}{4096n^4},
\end{aligned}$$

where in the second inequality we used Thm. 5 and in the third inequality we used Jensen's inequality with respect to the convex function $\mathbf{x}\mapsto|\mathbf{x}|^5$. □

4 Efficiently Approximating Functions with Small Representations

In this section, we show that it is possible to approximate a wide family of functions using ReLU neural networks, where the error decays exponentially with the depth. Specifically, we consider functions which can be approximated using a moderate number of multiplications and additions, where the values of intermediate computations are bounded (for example, a special case is any function approximable by a moderately-sized Boolean circuit, or a polynomial).

The key result to show this is the following, which shows that the multiplication of two (bounded-size) numbers can be approximated by a neural network, with error decaying exponentially with depth:

Theorem 6. *Let $f:[-M,M]^2\to\mathbb{R}$, $f(x,y)=x\cdot y$ and let $\epsilon>0$ be arbitrary. Then exists a ReLU neural network g of width $4\left\lceil\log\left(\frac{M}{\epsilon}\right)\right\rceil+13$ and depth $\left\lceil 2\log\left(\frac{M}{\epsilon}\right)\right\rceil+9$ satisfying*

$$\sup_{(x,y)\in[-M,M]^2}|f(x,y)-g(x,y)|\leq\epsilon.$$

The idea of the construction is that depth allows us to compute highly-oscillating functions, which can extract high-order bits from the binary representation of the inputs. Given these bits, one can compute the product by a procedure resembling long multiplication:

Proof of Thm. 6. We begin by observing that by using a simple linear change of variables on x, we may assume without loss of generality that $x\in[0,1]$, as we can just rescale x to the interval $[0,1]$, and then map it back to its original domain $[-M,M]$, where the error will multiply by a factor of $2M$. Then by requiring accuracy $\frac{\epsilon}{2M}$ instead of ϵ, the result will follow.

The key behind the proof is that performing bit-wise operations on the first k bits of $x \in [0,1]$ yields an estimation of the product to accuracy $2^{1-k}M$. Let $x = \sum_{i=1}^{\infty} 2^{-i}x_i$ be the binary representation of x where x_i is the i^{th} bit of x, then

$$\begin{aligned} x \cdot y &= \sum_{i=1}^{\infty} 2^{-i}x_i \cdot y \\ &= \sum_{i=1}^{k} 2^{-i}x_i \cdot y + \sum_{i=k+1}^{\infty} 2^{-i}x_i \cdot y. \end{aligned} \tag{12}$$

But since

$$\left| \sum_{i=k+1}^{\infty} 2^{-i}x_i \cdot y \right| \leq \left| \sum_{i=k+1}^{\infty} 2^{-i} \cdot y \right| = 2^{-k} |y| \leq 2^{1-k}M,$$

Eq. (12) implies

$$\left| x \cdot y - \sum_{i=1}^{k} 2^{-i}x_i \cdot y \right| \leq 2^{1-k}M.$$

Requiring that $2^{2-k}M \leq \frac{\epsilon}{2M}$, it suffices to show the existence of a network which approximates the function $\sum_{i=1}^{k} 2^{-i}x_i \cdot y$ to accuracy $\frac{\epsilon}{2}$, where $k = 2\left\lceil \log\left(\frac{8M}{\epsilon}\right)\right\rceil$. This way both approximation will be at most $\frac{\epsilon}{2}$, resulting in the desired accuracy of ϵ.

Before specifying the architecture which extracts the i^{th} bit of x, we first describe the last 2 layers of the network. Let the penultimate layer comprise of k neurons, each receiving both y and x_i as input, and having the set of weights $\left(2^{-i}, 1, -1\right)$. Thus, the output of the i^{th} neuron in the penultimate layer is

$$\left[2^{-i}y + x_i - 1\right]_+ = 2^{-i}x_i y.$$

Let the final single output neuron have the set of weights $(1, \ldots, 1) \in \mathbb{R}^k$, this way, the output of the network will be $\sum_{i=1}^{k} 2^{-i}x_i \cdot y$ as required.

We now specify the architecture which extracts the first most significant k bits of x. In Telgarsky [2016], the author demonstrates how the composition of the function

$$\varphi(x) = [2x]_+ - [4x - 2]_+$$

with itself i times, φ^i, yields a highly oscillatory triangle wave function in the domain $[0,1]$. Furthermore, we observe that $\varphi(x) = 0 \; \forall x \leq 0$, and thus $\varphi^i(x) = 0 \; \forall x \leq 0$. Now, a linear shift of the input of φ^i by 2^{-i-1}, and composing the output with

$$\sigma_\delta(x) = \left[\frac{1}{2\delta} - \frac{1}{4\delta} + \frac{1}{2}\right]_+ - \left[\frac{1}{2\delta} - \frac{1}{4\delta} - \frac{1}{2}\right]_+,$$

which converges to $\mathbb{1}_{[x \geq 0.5]}(x)$ as $\delta \to 0$, results in an approximation of $x \mapsto x_i$:

$$\sigma_\delta\left(\varphi^i\left(x - 2^{-i-1}\right)\right).$$

We stress that choosing δ such that the network approximates the bit-wise product to accuracy $\frac{\epsilon}{2}$ will require δ to be of magnitude $\frac{1}{\epsilon}$, but this poses no problem as representing such a number requires $\log\left(\frac{1}{\epsilon}\right)$ bits, which is also the magnitude of the size of the network, as suggested by the following analysis.

Next, we compute the size of the network required to implement the above approximation. To compute φ only two neurons are required, therefore φ^i can be computed using i layers with 2 neurons in each, and finally composing this with σ_δ requires a subsequent layer with 2 more neurons. To implement the i^{th} bit extractor we therefore require a network of size $2 \times (i+1)$. Using dummy neurons to propagate the i^{th} bit for $i < k$, the architecture extracting the k most significant bits of x will be of size $2k \times (k+1)$. Adding the final component performing the multiplication estimation will require 2 more layers of width k and 1 respectively, and an increase of the width by 1 to propagate y to the penultimate layer, resulting in a network of size $(2k+1) \times (k+1)$. □

The previous result shows that multiplication can be performed accurately by deep networks. Moreover, additions can be computed by ReLU networks exactly, using only a single layer with 4 neurons: Let $\alpha, \beta \in \mathbb{R}$ be arbitrary, then $(x, y) \mapsto \alpha \cdot x + \beta \cdot y$ is given in terms of ReLU summation by

$$\alpha [x]_+ - \alpha [-x]_+ + \beta [y]_+ - \beta [-y]_+ .$$

Thus, any function which can be approximated by a reasonable amount of operations involving additions and multiplications, can also be approximated well by moderately-sized networks. This is formalized in the following theorem:

Theorem 7. *Suppose $\mathcal{F}_{t,M,\epsilon}$ is the family of functions on the domain $[0,1]^d$ with the property that $f \in \mathcal{F}_{t,M,\epsilon}$ is approximable to accuracy ϵ with respect to the infinity norm, using at most t operations involving weighted addition, $(x, y) \mapsto \alpha \cdot x + \beta \cdot y$, where $\alpha, \beta \in \mathbb{R}$ are fixed; and multiplication, $(x, y) \mapsto x \cdot y$, where each intermediate computation stage is bounded in the interval $[-M, M]$. Then there exists a universal constant c, and a ReLU network g of width and depth at most $c\left(t \log\left(\frac{1}{\epsilon}\right) + t^2 \log(M)\right)$, such that*

$$\sup_{\mathbf{x} \in [0,1]^d} |f(\mathbf{x}) - g(\mathbf{x})| \leq 2\epsilon.$$

Proof. We begin by monitoring the rate of growth of the error when performing either an addition or a multiplication. Suppose that the given input $\tilde{a}, \tilde{b}$ is of distance at most $\delta > 0$ from the desired target values a, b, i.e., $|a - \tilde{a}| \leq \delta, \left|b - \tilde{b}\right| \leq \delta$. Then for addition we have

$$\left|(a+b) - \left(\tilde{a} + \tilde{b}\right)\right| \leq |a - \tilde{a}| + \left|b - \tilde{b}\right| \leq 2\delta,$$

and for multiplication we compute the product error estimation

$$\begin{aligned} \left|\tilde{a} \cdot \tilde{b} - a \cdot b\right| &\leq |(a+\delta) \cdot (b+\delta) - a \cdot b| \\ &= \left|\delta(a+b) + \delta^2\right|. \end{aligned}$$

Now, on top of the cumulative error, we add the error of approximating the product of $a \cdot b$, which we may assume is at most δ for each product computed in the series of operations performed, if enough bits are used as in the proof of Thm. 6, to have that

$$\left|\delta(a+b) + \delta^2 + \delta\right| \leq 3M\delta.$$

That is, at each stage the error grows by at most a multiplicative factor of $3M$. After t operations, and with an initial estimation error of δ, we have that the error is bounded by $(3M)^{t-1}\delta$. Choosing $\delta \leq (3M)^{1-t}\epsilon$ to guarantee approximation ϵ, we have from Thm. 6 that each operation will require at most

$$4\left\lceil \log\left(\frac{M(3M)^{t-1}}{\epsilon}\right)\right\rceil + 13 \leq c\left(\log\left(\frac{1}{\epsilon}\right) + t\log(M)\right)$$

width and

$$2\left\lceil \log\left(\frac{M\left(3M\right)^{t-1}}{\epsilon}\right)\right\rceil + 9 \leq c\left(\log\left(\frac{1}{\epsilon}\right) + t\log\left(M\right)\right)$$

depth for some universal $c > 0$. Composing the networks performing each operation, we arrive at a total network width and depth of at most

$$c\left(t\log\left(\frac{1}{\epsilon}\right) + t^2\log\left(M\right)\right).$$

Now, our target function is approximated to accuracy ϵ by a function which our network approximates to the same accuracy ϵ, for a total approximation error of the target function by our network of 2ϵ. □

5 The Formal Separation Result

With the previous results in hand, we are now finally ready to present our main result, which formally shows how depth can be exponentially more valuable than width as a function of the target accuracy ϵ:

Corollary 3. *Suppose $f \in C^2 \cap \mathcal{F}_{t(\epsilon),M(\epsilon),\epsilon}$, where $t\left(\epsilon\right) = \mathcal{O}\left(\text{poly}\left(\log\left(\frac{1}{\epsilon}\right)\right)\right)$ and $M\left(\epsilon\right) = \mathcal{O}\left(\text{poly}\left(\frac{1}{\epsilon}\right)\right)$. Then approximating f to accuracy ϵ in the L_2 norm using a fixed depth ReLU network requires width at least $\frac{c}{\epsilon}$, whereas there exists a ReLU network of depth and width at most $p\left(\log\left(\frac{1}{\epsilon}\right)\right)$ which approximates f to accuracy ϵ in the infinity norm, where c, p are a constant and a polynomial respectively, depending solely on f.*

The proof of Corollary 3 follows immediately from Corollary 2 and Thm. 7. While the lower bound is straightforward, for the upper bound observe that Thm. 7 implies an ϵ approximation by a network of width and depth at most

$$c\left(t\left(\frac{\epsilon}{2}\right)\log\left(\frac{2}{\epsilon}\right) + \left(t\left(\frac{\epsilon}{2}\right)\right)^2\log\left(M\left(\frac{\epsilon}{2}\right)\right)\right),$$

which by the assumption of Corollary 3, can be bounded by

$$p\left(\log\left(\frac{1}{\epsilon}\right)\right)$$

for some polynomial p which depends solely on f.

References

Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *29th Annual Conference on Learning Theory*, pages 698–728, 2016.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

Ronen Eldan and Ohad Shamir. The power of depth for feedforward neural networks. In *29th Annual Conference on Learning Theory*, pages 907–940, 2016.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.

Matus Telgarsky. benefits of depth in neural networks. In *29th Annual Conference on Learning Theory*, pages 1517–1539, 2016.